



Data Engineering for Smart City Traffic Prediction

Ghatoth mishra

Independent Researcher

ghatothmishra@gmail.com

Abstract

Traffic flow prediction is a crucial component of developing intelligent transportation systems in smart cities. The core goal is to investigate the data engineering approaches and methodologies that support traffic flow prediction. Structured but diverse traffic flow data sets generated by sensor networks and other sources are ingested, processed, and stored in data platforms that enable operational real-time or near-real-time predictions and alternative exploratory offline analyses. Effective data engineering process solutions for traffic flow prediction require the careful design of all stages—from data sources to model application—because the dynamic and still little-understood nature of traffic flow data makes dedicated traffic prediction models sensitive to a variety of factors, including data quality, time frame, spatial resolution, external data, feature representation, and algorithm selection. Consideration of these factors can yield appropriate solutions across data scenario alternatives: data from a limited number of sensors, real-time prediction with external data to better represent special events, online learning for concept drift adaptation, and feature engineering with new perspectives or resources.

Traffic flow prediction represents an operational application in the complex and multidisciplinary scenarios of a smart city. Intelligent transport systems rely on timely and accurate real-time predictions to optimize vehicle distribution, reduce waiting times, increase passenger satisfaction, and enable vehicle tracking. Such predictions also support external decision-making processes that require support from an intelligent system or subsystem. Usually classified as time-series forecasting, traffic flow prediction aims to inferring future values of a time-ordered series generated by one or more object through a range sensor, such as microwave, loop, infrared, or video camera sensor. The tasks for traffic flow prediction comprises sensor networks and the incoming traffic flow data streams, as well as multiple connected external data source that contribute to broaden the representation of predicted-related phenomena—e.g. weather conditions, official event schedule—and support data fusion.

Keywords : Traffic flow prediction, Smart cities, Data engineering, Spatio-temporal data, Intelligent transportation systems (ITS), Big data analytics, Internet of Things (IoT), Real-time data processing, Data pipelines, Feature engineering, Graph neural networks (GNN), Deep learning, Time series forecasting, Data fusion (multi-source integration), Edge computing.

1. Introduction

Smart cities' traffic systems produce and consume vast amounts of data in real time. These data can be exploited to forecast future traffic conditions and to manage and operate the entire traffic ecosystem more efficiently and intelligently. Predicting traffic is complex due to traffic dynamics, road conditions, and surrounding environmental

effects. Data engineering covers all aspects of data handling, focusing on data scientists and machine-learning engineers' requirements. For traffic prediction, smart cities' data landscape should first be addressed. Prediction dashboards involve sensor networks generating continuous streams of traffic data. Other internal sources such as traffic events, weather conditions, or road construction are often overlaid. Furthermore, open datasets provide a historical context



supporting anomalies and model evaluation. After ingesting and integrating data from various sources, data management is tackled. Time-series databases and data lakes are the most common data repositories. Data-model/schema design is considered, depending on the model training and serving needs.

Traffic prediction typically requires the incorporation of external factors as features: features that account for the temporal and spatial components of the traffic phenomenon in a timely manner. When exploring advanced methods, traffic-flow networks can also be leveraged to represent the problem in spatial terms, using the input–output architecture of flow networks. Modeling approaches mainly include traditional statistical models, autoregressive models, and multiple-variable regression, along with deep-learning models, which are widely used. When aiming for real-time prediction, streaming analytics techniques are applied, with low-latency predictions deployed in production. In this case, online learning represents and copes with the dynamics of traffic-flow behaviors by frequently updating models with new data.

1.1. Background and Significance

Traffic congestion and related environmental issues in smart cities exert a direct strain on residents' quality of life. Short-term traffic flow prediction enables formal operators and end users to be better prepared for anticipated travel conditions. The prediction task requires sophisticated data engineering approaches to assemble proper feature sets from the increasingly rich and diverse urban data landscape. These include an expansive set of sensor networks and data archives, which must be shaped into a suitable model training and inference environment. In most existing work, a prediction topology is developed but not the entire data and processing pipeline. This section fills that gap. A small set of generic real-time prediction and operational aspects, such as online learning and adaptation in low-latency use cases, are also outlined.

Communications Technology and IoT Data Streams Traffic flow prediction relies heavily on fixed-position sensors,

particularly inductive loop detectors. These sensor networks, particularly when combined with traffic cameras, are supplemented by other traffic sources—such as GPS traces from taxis and ride-sharing services, social media activity, and identified traffic accidents—that are contributing to a richer and more diverse traffic ecosystem in many smart cities. Nonetheless, movements related to long-distance traffic, tourism, and logistics remain underrepresented. Recent research has also highlighted the value of integrating local weather data in traffic prediction models for urban mobility.




Fig 1: Optimizing Traffic Flow in Smart Cities


1.2. Research design

Research on data engineering approaches for traffic flow prediction in smart cities is guided by three concrete objectives. The first objective is to identify relevant data in the smart city traffic ecosystem—the potential role of any data, from traffic sensors, open data, or third-party sources, is considered. For this reason, data-related questions must be answered: How is data collected? Where is it stored? How is it made available for traffic prediction? Data engineering



answers these questions by describing the steps from data ingestion to operational model development.

The data landscape, data acquisition, and storage patterns of smart city traffic infrastructure are examined. Traffic prediction is then viewed from the lens of supervised machine-learning scenarios, which comprise a training and inference stage. The design of these stages is a key step tightly linked to the data. Several strategic steps combine to form the research design: the feature space—the set of characteristics employed to forecast future traffic—is defined, and methods for dynamic prediction at a low latency are surveyed.



2. Data Landscape in Smart City Traffic Systems

The data landscape of traffic systems is heterogeneous, composed of data partitions cycling through different vectors of volume and velocity. Deployed sensor networks produce continuous streams of data, which can be consumed through time-series databases or distributed processing systems for low-latency analytical applications. City-wide open-data initiatives provide supplementary sources of information from government agencies, traffic reports, weather stations, point-of-interest data, and social media activity.

All data types are consumed to create feature sets used for predictive modeling of traffic flows and congestions at various time scales. The production of rich feature sets from the integrated data landscape is a data engineer task, including data mining, feature generation and selection, and feature representation. Part of the feature set is transitively defined using a spatial database schema able to represent the city as a directed graph without loops or crosses. Weighting features associated to network segments and nodes can be easily generated as a side effect of monitoring or data-mining operations, enabling their integration to real-state anomaly detection and prediction tasks.

Equation 1: AR, MA, ARMA, ARIMA, SARIMA derivations (step-by-step)

Assume flow depends linearly on the last p values:

$$y_t = c + \phi_1 y_{t-1} + \dots + \phi_p y_{t-p} + \varepsilon_t$$

where ε_t is zero-mean noise.

Derivation from “linear regression on lag features”:

1. Build features: $\mathbf{x}_t = [1, y_{t-1}, \dots, y_{t-p}]^T$
2. Parameters: $\boldsymbol{\beta} = [c, \phi_1, \dots, \phi_p]^T$
3. Then $y_t = \mathbf{x}_t^T \boldsymbol{\beta} + \varepsilon_t$

Instead of lags of y , model lags of error:

$$y_t = \mu + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \dots + \theta_q \varepsilon_{t-q}$$

Combine AR and MA:

$$y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \varepsilon_t + \sum_{j=1}^q \theta_j \varepsilon_{t-j}$$

Traffic data is often non-stationary (trend/seasonality), so difference it.

First difference:



$$\Delta y_t = y_t - y_{t-1}$$

Second difference:

$$\Delta^2 y_t = \Delta y_t - \Delta y_{t-1}$$

In general, with operator $(1-B)$ where $By_t = y_{t-1}$:

$$\Delta y_t = (1-B)y_t, \Delta^d y_t = (1-B)^d y_t$$

ARIMA definition:

Apply ARMA to the differenced series $w_t = \Delta^d y_t$:

$$w_t = c + \sum_{i=1}^p \phi_i w_{t-i} + \varepsilon_t + \sum_{j=1}^q \theta_j \varepsilon_{t-j}$$

This corresponds to the article's mention of **ARIMA/SARIMA** as classical statistical methods for traffic forecasting.

Data Engineering Approaches for...

If there is seasonality with period s (e.g., daily pattern), add seasonal differencing:

$$\Delta_s y_t = y_t - y_{t-s}$$

and combine both:

$$w_t = (1-B)^d (1-B^s)^D y_t$$

3. Data Ingestion and Integration

Data ingestion and integration represent essential pillars in the data engineering landscape of smart city traffic systems. Data consumption begins with sensor networks, whose dedicated protocols deploy low-cost, low-power, low-bandwidth radio transmitters to minimize communication costs for devices with limited capacity. Traffic volume detection by inductive loops or magnetometers serves as the

primary input. Mobile and GPS sensor data from taxis, public transport, and private vehicles equip traffic prediction with richer datasets on travel paths and origin–destination flows. Cameras detect fine-grained vehicle class distributions, while WLAN probes deliver people movement patterns.

Virtual sensing compresses traffic measurements by airborne radio signals, helping major monitoring stations to predict short-term flows on adjacent links. Thermal infrared and short-term radar detect probabilities of crosswalk occupancy and pedestrian–vehicle interactions. Integration with external data sources, both internal and open, further enriches real-time predictions. Weather forecasts shape anticipated changes in travel behavior; event calendars revolve traffic volume and road use patterns. The city's cloud infrastructure federates data storage across its municipal departments while opening historical datasets for external exploitation. Urban systems of greatest interest for research and innovation are released as open data on an external portal. Data are prepared and cleaned offline before becoming accessible, and the city's initiatives apply for scientific investigator authorization in triggering specific elements of the data environment.

3.1. Sensor Networks and IoT Data Streams

Data from sensor networks and real-time IoT streams enables the prediction of future traffic conditions based on already obtained similar temporally-spatial traffic conditions or data from the roadway segments with parallel condition. In addition to conventional methods, deep learning approaches also have grown in popularity. The models are designed to build deep neural networks from the spatio-temporal attributed graph data that represent the traffic system. For prediction, the models explore both historical and real-time data, without consideration of the traffic lights or intersections. The appearance of the traffic lights and intersections may lead to traffic jams, while the change of saturation flow may also affect future traffic conditions. Therefore, research finds effective methods by adding special attention designs.

Research also finds ways to predict traffic flow from the history and spatial attribute perspective, without the need for an external spatial-temporal sequence model. In transportation, prediction function can be defined by transportation theory. For prediction, one may consider all the roadway segments in the city as a connected graph. Since all the road segments are connected, the traffic flow can be transferred around. Convolution methods in the graph-based model extract the local interrelations while the spatio-temporal convolution extracts the data in the temporal domain from historical records.




Fig 2: Integration and Exploitation of Sensor Data in Smart Cities

3.2. External Data Sources and Open Data

External events, the natural environment, and demographic trends have been shown to influence traffic flow. Traffic flow prediction is therefore regularly augmented with additional data and information from external sources. The wide range of available datasets from different external providers can assist attempts to integrate external information into traffic flow forecasting. Providing information about weather influences is particularly widespread; for example, meteorological conditions affect traffic flows due to changes in road conditions (ice, snow, etc.). Additionally, the combination of open public data with

sensor data generally improves traffic flow prediction. Other data sources used in combination with historical sensor data include temporal information (day of the week, month, year) and public holiday information.

During the era of big data, urban issues have drawn more attention, which has led to diverse geographical and thematic open datasets under the “open data” concept. Furthermore, the increase in the number of datasets has made it challenging to shortlist appropriate auxiliary datasets that can assist in data-based applications, such as traffic flow prediction. Open datasets, other than traffic flow historical data, can accommodate long-term prediction processes. Open datasets related to land use types, spatial populations, public transport, social media, accidents, and points of interest are important for urban analytics in smart cities. The seamless integration of multiple auxiliary datasets concentrated on meeting the requirements of various urban issues can enhance traffic flow prediction capability. Available social media data, especially Twitter data related to traffic transportation, can reflect travel intentions and help to improve prediction results. Weather condition datasets, which reflect the road environment and vehicle operation conditions, are also essential sources for a wide range of practical applications.

4. Data Storage and Management

Time-series data streaming into smart city traffic systems are characterized by significant volumes along several dimensions, including spatial, temporal, and data types. Traffic-flow predictions are influenced not only by historical data from the area surrounding the target location but also by information from adjacent areas and other modes of transportation. Traffic-flow prediction tasks require special storage systems that facilitate high-velocity writing and, at the same time, offer low-latency reading capabilities. Time-series databases, capable of handling several million write operations per second, are often used for the ingestion of sensor data, while storage data lakes that leverage a cloud storage service (e.g., S3) are employed for the storage of



historical and non-time-series data. Data-analytics models commonly require complex architectural designs. Signals need to be organized, cleaned, and indexed for efficient joint access, typically orchestrated with the help of an Apache Airflow or equivalent workflow orchestrator. Supporting services for online prediction—e.g., warning users of an accident or an important event—or for systematically managing existing models are built using serverless architectures that provide easy scaling and reduce operational management.

Data stored in a data lake require different treatments. Data types can be heterogeneous; therefore, a tailored schema-on-read is adopted. For temporal regression tasks, the influence of time on the output signal must be taken into account. In addition to representing time in the model, other spatial and temporal features can be added to the data preparation pipeline. Other sources of data can implicitly introduce features in the machine-learning model. Roads can be represented as graphs, where nodes represent the crossings of all roads used for signalization and the edges have a label that indicates the label of the type of road (highway, street, etc.), or at least its width. Distances among nodes can be computed and added to the dataset.

Equation 2: Multi-sensor, multi-output (network) forecasting

Let there be N sensors (or road segments). Define:

$$\mathbf{y}_t \in \mathbb{R}^N$$

where $[\mathbf{y}_t]_i$ is sensor i 's flow at time t .

With lookback p , define the spatio-temporal input tensor:

$$\mathbf{X}_t = [\mathbf{y}_t, \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-p+1}] \in \mathbb{R}^{N \times p}$$

Forecast:

$$\hat{\mathbf{y}}_{t+h} = f(\mathbf{X}_t)$$

Flatten $\mathbf{X}_t$ into a vector $\mathbf{x}_t \in \mathbb{R}^{Np}$, then:


$$\hat{\mathbf{y}}_{t+h} = \mathbf{W}\mathbf{x}_t + \mathbf{b}$$

where $\mathbf{W} \in \mathbb{R}^{N \times Np}$.

4.1. Time-Series Databases and Data Lakes

Traffic prediction in smart cities requires an explosion of data, sometimes of different types and at different velocities, and thus different architectures are needed not only for data ingestion and preprocessing but also for data storage and management. For example, a traditional relational database may be not sufficient, given that traffic data is mainly time series. A time-series database (TSDB) may be more appropriate for such data because it is useful for monitoring time-series data and optimizing the operations involved in that process. A TSDB is designed to handle metric data that is indexed by time (temperature, humidity, speed, for example) that must be retrieved and grouped by time intervals. Optimization can also be performed on the database size (given that strings cannot be used) as well as an aggregation on the timestamps, allowing faster data retrieval, minimization of the index overhead, and reduced storage footprint. Some TSDB solutions support wider data types, such as the object-relational database solution TimescaleDB for PostgreSQL-based systems. Additionally, time-series databases such as InfluxDB are not limited to time-series data; they are also being adopted in other applications.

A data lake may also be useful for a smart city traffic system, enabling storage and retrieval of a larger variety of data sources, such as videos and images captured with a medium or low capture rate. A data lake supports the storage and retrieval of unstructured, structured, and semi-structured data with different storage capacities and velocities across its different layers, for example, a hot or cold storage layer. A data lake does not impose a predefined structure or schema, thus allowing data to be stored without preorganization. Hence, it can accommodate new data types as needed.



4.2. Schema Design and metadata Management

Time-Series Databases (TSDB) preserve a subset of the sensor measurements and are particularly relevant for the storage of road occupancy data. The high granularity of these measurements makes the cost of storage a nontrivial aspect. A simple schema that enables pivoting along the time dimension optimizes the efficiency of the data population step, clustering values from different measures, but also from different sensors in the same time interval. The TSDB is complemented by a simple file-storage data structure to absorb other commonly used features during the prediction process, such as meteorological conditions and holidays.

Data lakes are well-suited for this traffic data ecosystem, as they natively support the storage of any kind of data and make it possible to organize, categorize, and label data using metadata, rather than fixing a schema a priori. Road mileage is diverse in terms of number of lanes and road type. Metadata describing a road segment can be used to extract specific features for certain bunches of similar stretches, removing noise and irrelevant data in the training and validations of ML models.

5. Feature Engineering for Traffic Prediction

Feature-engineering methodologies are tailored for drawing predictive models from training datasets across different labour strata. Spatial and temporal characteristics should be externalised into predictor variables, using the training dataset in conjunction with the graph structure of the sensor network to enhance inherent properties embedded in the data. Partitioning through separate geographical detection zones (possibly centred on junctions) and the instantaneous flow from neighbouring zones can also be valuable predictors, particularly in statistical methods. Thereby, all such variables thus generated must be endowed with metadata (geographical range of effect, granularity, physical dimension, temporal description), ensuring that laboratory personnel will comprehend the attributes during cross-correlation, representation design and/or modelling.

Furthermore, such features must be delivered to the main computer for experimentation. Because distance features are usually redundant, the search for the best subset of predictors should employ progressive and/or chaotic forward selection. Until now, feature selection and model calibration were static operations, but sensors are located in evolving non-congruent preferential paths. To capture temporal localities, the training set should be formed by subsets of earlier measurements, keeping within the arrest function of traffic states for low-minimum-measure-density flows. The need for traffic-flow prediction extends further into time, since operation-response anticipations derive long-short cycle attribute functions. Consequently, minimising prediction flow–error covariance for each request localises the pattern–data match.

5.1. Spatial and Temporal Features

Incorporating features capturing the spatial and temporal dimensions of traffic flow is a common approach in traffic prediction. Research literature suggests two types of temporal features: categorical features which model daily and weekly seasonality (e.g., day of the week), and continuous features which model longer patterns (e.g., time of day, or Time Since Midnight). Seasonal decompositions, which can also be learned by ML-based models like LSTM, have also been proposed in the literature.



In addition to time of day, day of week and public holidays, time of year could also be included as a potential predictor. Bank holidays, school holidays and other such events can have strong influences upon system demand. A joint bank-holiday representation by multiple countries could also be used in modelling airport flows.

On the spatial side, how different locations relate to each other plays a key role in the prediction process. Relationships of nearby locations are often represented in a pairwise manner (e.g. in an N-by-N adjacency matrix). The Global Spatial Features Scaling (GSFS) technique has shown significant predictive additive value by modelling such pairwise spatial influences for Scenes 2, 5 and 6 in the SFO dataset.

Finally the provision of spatial-temporal attention mechanisms has provided predictive gains in capturing more complex spatial-temporal dynamics when resources allow, for example using Wang et al.'s Spatio-Temporal Attention Convolutional Recurrent Neural Network (STACRNN) model.




Fig 3: Temporal-Spatial Traffic Flow Prediction Model

5.2. Graph-Based Representations

Graph-based representations of traffic infrastructure enable more complex interactions between spatial and temporal predictors, exploiting the structure of the road network. External sources such as weather or calendar events, which require specific spatiotemporal patterns to be useful for learning, are naturally integrated within the representation.

The road network is modeled as a graph, with intersections as nodes and road segments as edges; directed edges account for one-way roads and help in defining spatial neighborhoods for a traffic flow time series. The choice of functional representation of the system influences the learning problem: if a traffic-dependent decreasing flow function in edges is considered, the predicted structure reveals hot spots of traffic congestion.

Road traffic is a self-driven process expressing local behavior but communicating flows on the intersection. The self-driven process allows predicting the future entry and exit from an intersection using only its own past flow. The incoming and outgoing fluxes of the four directions are conditioned on the four directions of the other intersections. Predictive models for each direction can express structural synopses directly related to congestion processes. Flow conservation can impose constraints directly in the learning process for each occupied edge. Probabilities can be used to emulate unknown effects conditioned on the entering intersections. Dynamical Expanding Graphs generated by the traffic events describe aspects related to traffic congestion. The classification associated with bottlenecks can help to understand a traffic jam event.

6. Modeling Approaches

Based on the nature of the data, its ingestion, integration, storage, and feature engineering strategies, a variety of statistical, machine learning, and deep learning approaches have been applied for short-term traffic flow forecasting at single locations, predicting multiple locations, and predicting the future traffic condition for the whole network or subnetwork. An additional distinction can be made between dedicated time-series traffic predictions and general machine learning tasks where traffic flow is considered a feature among others, such as accident detection and short-term prediction, traffic condition classification, travel time prediction, etc.



Statistical approaches such as autoregressive integrated moving average (ARIMA) have been extensively used in traffic flow forecasting. However, ARIMA forecasting model and its variations are only suitable for single-location time-series forecasting because they do not consider the correlation of traffic conditions at adjacent locations. Classical prediction models such as multivariate linear regression and support vector regression can be easily extended to consider multiple locations. More advanced models, such as neural networks (NNs), Gaussian process regression (GPR), k-nearest neighbor, multilayer perceptrons (MLPs), and long short-term memory (LSTM) neural networks, are used for both single-location and multiple-location predictions. Unlike purely statistical models, these methods allow the automatic extraction of suitable prediction features from historical data and additionally support nonlinear relationships. However, the prediction performance is usually limited by the generalization ability of the model. Thus, it is inevitable that not only past data are used but also advanced artificial intelligence (AI) methods are applied for real-time traffic prediction — for both single-location prediction and prediction of multiple major locations in a city.

Equation 3: Graph-based representation and GCN derivation (step-by-step)

Let:

- $A \in \mathbb{R}^{N \times N}$: adjacency matrix ($A_{ij} = 1$ if connected, else 0)
- D : degree matrix, $D_{ii} = \sum_j A_{ij}$

Graph Laplacian:

$$L = D - A$$

Raw Amakes sums depend on node degree. Normalize:

$$\tilde{A} = A + I, \tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$$

Symmetric normalized adjacency:

$$\hat{A} = \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}}$$

Let node features be $H^{(0)} = X \in \mathbb{R}^{N \times F}$.

One GCN layer:

$$H^{(1)} = \sigma(\hat{A}H^{(0)}W^{(0)})$$

- $\hat{A}H^{(0)}$: aggregates neighbor information
- $W^{(0)}$: learnable linear map
- $\sigma(\cdot)$: nonlinearity (ReLU, etc.)

Step-by-step meaning:

1. Add self-loops: $\tilde{A} = A + I$
2. Compute degrees: $\tilde{D}$
3. Normalize: $\hat{A} = \tilde{D}^{-1/2} \tilde{A} \tilde{D}^{-1/2}$
4. Neighbor aggregation: $M = \hat{A}H^{(0)}$
5. Feature mixing: $MW^{(0)}$
6. Nonlinearity: $\sigma(\cdot)$

6.1. Statistical and Classical Methods

Traffic flow prediction can be tackled with statistical and classical time-series forecasting techniques, such as ARIMA or SARIMA. These models exploit temporal dependencies within the data and have successfully provided short-term predictions of traffic flows. However, when dealing with large-scale systems (i.e., models trained across multiple sensors), learning a sufficient number of temporal patterns for all locations becomes difficult due to the limited number of training samples.

Regression approaches using linear models (such as regression trees or lasso regression) have been applied to the traffic flow pre-diction problem, determining the influence



of temporal and traffic features on different locations and enabling the identification of interesting patterns. Additionally, traditional models such as support vector regression and Bayesian regression have been applied to individual sensors, capitalizing on the sensors' abilities to model non-linearity and to auto-adapt to traffic flow. Even if these approaches address the non-stationary behavior of traffic, they still do not properly model the complex spatial dependencies of the underlying system.

In recent years, a number of approaches have exploited these dependencies with great success. The idea is to represent the system as a traffic graph described by a set of nodes connected based on spatial proximity and traffic relationships. In this case, the aim of prediction is to model the intensity of the flow (volume and/or speed) for each leg during a fixed horizon. Most approaches use the historical traffic data associated with single nodes or whole roads as features and compute traffic volume, speed and time-lagged features to capture the temporal relationship between the connections.

6.2. Machine Learning for Traffic Flow

Recent advances in Artificial Intelligence have shown that Machine Learning (ML) in general and Deep Learning (DL) in particular belong to the state-of-the-art methods in data science in general. Such new methodologies have been successfully applied also on traffic flow prediction tasks on benchmarks using static data.

Experiments have shown that the choice of ML algorithms and of their parameters plays a crucial role for the performance of the prediction task, thus favoring the actualization of an intensive search of the best algorithm/parameters pair. Nevertheless, in real time tasks it is not always possible to apply any of the mentioned stages, especially that of hyper-parameters tuning, or using ML for the prediction task at all. Manifold Traffic Flow Prediction (MTFP) is a very challenging task in which traffic flow conditions of various sensors are predicted, using the past behaviour of many sensors as input. The prediction is heavily attacked by occurrences like accidents, traffic jams

or climate adversities. The concrete task is considered as a multi-output prediction problem in which the flow of $N+1$ sensors within a region is predicted during a time period Δ equal to 30 minutes, using as input data that collected during the last Δ minutes. The meta-ensemble of local predictors easily outperforms the results of a single (global) ML model.

7. Real-Time Prediction and Operational Systems

Operational systems require real-time traffic flows and predictions throughout the operational horizon. Time-critical processes, such as adaptive traffic light management and rerouting recommendations, require very low-latency responses. Such processes should ideally be realized as streaming systems, with data flow from sensor networks through prediction and control directly into action in the city. Stream processing introduces new challenges for model design. Heavy-tailed response time distributions and frequent updates are inherent to these processes. Models must meet very stringent latency requirements, where even a few seconds can create undesired effects. Retaining light configurations for a few seconds can lead to queues forming at the upstream corner.

Modeling approaches must also consider these constraints. Data-driven models can realize very complex behavior but require substantial amounts of data for training. A big portion of the prediction duration is often consumed by model loading. Lightweight and shallow models allow for very low latencies. Their performance is often tuned by ensemble approaches with a number of lightweight models. Implementing streaming analytics for very low-latency responses can also drive the use of online learning techniques for model adaptation. Insights from the modeling experiments need to be considered and prediction engines exposed.




Fig 4: Real-Time Prediction and Operational Systems

7.1. Streaming Analytics and Low-Latency Inference

Model development for traffic prediction might be carried out in a dedicated server or in the cloud and can be scheduled in a recurrent manner so as to retrain an inference pipeline, with a clear separation between preparation and inference. At execution time, features are computed, points of inference are selected, and parallel predictions are triggered for all the required models and pipelines. In some cases it may be desirable to have low-latency predictions that leverage real-time detected traffic conditions to feed ahead the predicted traffic state for road networks, temporarily propagated to the pure model as additional inputs. This can be implemented using low-latency stream-processing analytics infrastructures and is particularly useful for advanced driving systems, where information must be ready for very short time horizons.

Real-time traffic prediction methods exploit sensor density and focused workload to be performed only close to the prediction time. The problem of predicting flows in the immediate future (next e.g. 30 min) is tackled, and the generated traffic state is intended to be the initial condition of prediction models working into larger time horizons (e.g. 1–12 h ahead). The procedure relies on small neural networks predicting all the required flows at every time step and is now being boosted by tracking individual vehicles with an enhanced algorithm. The results from the prediction system are also the first step for another advanced

behaviour; a reduced set of traffic conditions may be used to feed trajectory planning algorithms for autonomous driving in the travels of EVs connected with the city information system.

7.2. Online Learning and Model Adaptation

Changes in traffic description and pattern, such as those from holidays, sporting events, or other macro-level phenomena (e.g., day-of-week and time-of-day), can render a model obsolete. Similarly, changes in the environment, such as roadwork, accidents, or natural phenomena, can also cause such issues. The parameters of models frequently need fine-tuning drop, especially in statistical models. Changes in traffic density may alter the accuracy of a regression model. Furthermore, because traffic from a specific location and time relies primarily on traffic a short time ahead, the relation becomes less strong as the forecasting distance increases.

Decision trees for traffic flow prediction can also exhibit an inclination subject to abrupt changes in the traffic flow of a single location, e.g. an accident. Such changes should be mitigated in real-time deployment. Models can be online learned continuously but only in hindsight, or their parameters can be updated periodically, only to a chosen extent, or with a moving window. One special case of online learning is semi-supervised learning, where only a small proportion of the data is labeled and supervised. Adaptation strategies that do not explicitly re-learn the model but take the prediction error into account may also be considered. Such strategies, often based on incremental learning, dynamic weighting or voting, assigning different confidence levels during aggregation, are possible for any model type. Even greater flexibility is offered by dynamic model selection.

8. Conclusion

Traffic prediction addresses one of the most classical problems in urban transport. It is important for numerous



applications, ranging from real-time navigation and routing systems to the anticipation of congestion and the management of urban infrastructures. During the last decade, the emergence of new sensor networks and the general trend towards the digitalization of urban environments have generated an unprecedented amount of data on traffic dynamics. All these data can provide significant improvements in prediction accuracy, as long as the complex spatio-temporal patterns embodied in traffic can be suitably captured.

In addition to the features offered by the sensor data, traffic flow prediction can also benefit from additional context-related or supporting variables, such as public transport status, civilian mobility preferences, weather conditions, and other events occurring in or near the city (soccer matches, concerts, etc.), that may introduce intermittence in the prediction process. Despite these opportunities, the exploited traffic flow data remain limited to few geographical sectors, neglecting the rest of the city. Novel spatio-temporal methodologies, based on graph-relational data representation and deep learning techniques, can ensure predictions across the entire road network using only local traffic flow data as input. Models address short-term traffic flow prediction (a few minutes ahead) for various time horizons (5, 15, and 30 minutes ahead) and for areas with different urban functions. The results demonstrate the effectiveness of the adopted methodologies and the importance of the spatio-temporal context in modeling traffic dynamics through data-driven approaches.

8.1. Future Trends

Future research on the data engineering side of smart city traffic flow prediction is likely to be influenced by four key trends:

- **Data ingestion and integration:** Traffic systems in smart cities may make use of increasingly comprehensive data resources, including data generated by dedicated sensor networks (e.g., for surveillance or environmental monitoring), by Internet-of-Things devices (e.g., air quality monitors distributed throughout the city) and by citizens

(e.g., data collections via mobile applications). Hence, developments such as continual streaming of time-series data from sensor networks and IoT deployments, more easy-to-use Application Programming Interfaces (APIs) for fetching data from external data sources and an increased availability of open data are crucial. Furthermore, the handling and management of these different types of data sources—some of them developed and maintained for the traffic system, others constituting external data sources that may still be useful for traffic prediction—are decisive for successful predictions and for their operational deployment in low-latency systems.

- **Predictive feature engineering:** The effective generation of informative features plays a vital role for traffic-prediction methods, both classical and machine learning. While statistical approaches require selection of the right relevant features, data- and model-driven methods can help. Spatial and temporal features, such as festival days or holiday seasons, improve time-series predictions and reduce prediction error. Moreover, representing traffic flow as a graph can increase the general applicability of models. Hence, new methods that help generate appropriate sets of features from existing or externally sourced data, as well as feature sets based on spatial-temporal concepts and relations, are expected to be fruitful.

9. References

1. Jia, T., & Yan, P. (2021). Predicting citywide road traffic flow using deep spatiotemporal neural networks. *IEEE Transactions on Intelligent Transportation Systems*, 22(5), 3101–3111.
2. Li, G., Knoop, V. L., & van Lint, H. (2021). Multistep traffic forecasting by dynamic graph convolution: Interpretations of real-time spatial correlations. *Transportation Research Part C: Emerging Technologies*, 128, 103185.



3. Li, J., Guo, F., Sivakumar, A., Dong, Y., & Krishnan, R. (2021). Transferability improvement in short-term traffic prediction using stacked LSTM network. *Transportation Research Part C: Emerging Technologies*, 124, 102977.
4. Inala, R. (2023). Big Data Architectures for Modernizing Customer Master Systems in Group Insurance and Retirement Planning. *Educational Administration: Theory and Practice*, 29(4), 5493-5505.
5. Manibardo, E. L., Laña, I., & del Ser, J. (2022). Deep learning for road traffic forecasting: Does it make a difference? *IEEE Transactions on Intelligent Transportation Systems*, 23(7), 6164–6188.
6. Wu, S. (2022). Spatiotemporal dynamic forecasting and analysis of regional traffic flow in urban road networks using deep learning convolutional neural network. *IEEE Transactions on Intelligent Transportation Systems*, 23(2), 1607–1615.
7. Ali, A., Zhu, Y., & Zakarya, M. (2022). Exploiting dynamic spatio-temporal graph convolutional neural networks for citywide traffic flows prediction. *Neural Networks*, 145, 233–247.
8. Lee, K., & Rhee, W. (2022). DDP-GCN: Multi-graph convolutional network for spatiotemporal traffic forecasting. *Transportation Research Part C: Emerging Technologies*, 134, 103466.
9. Afrin, T., & Yodo, N. (2022). A long short-term memory-based correlated traffic data prediction framework. *Knowledge-Based Systems*, 237, 107755.
10. Ren, Y., Jiang, H., Ji, N., & Yu, H. (2022). TBSM: A traffic burst-sensitive model for short-term prediction under special events. *Knowledge-Based Systems*, 240, 108120.
11. Wang, Y., Jing, C., & Guo, T. (2022). Attention based spatiotemporal graph attention networks for traffic flow forecasting. *Information Sciences*, 607, 869–883.
12. Wang, Q., Jiang, H., Qiu, M., & others. (2022). TGAE: Temporal graph autoencoder for travel forecasting. *IEEE Transactions on Intelligent Transportation Systems*.
13. Xue, R., Zhao, S., & Han, F. (2022). An embedding-driven multi-hop spatio-temporal attention network for traffic prediction. *IEEE Transactions on Intelligent Transportation Systems*.
14. Guo, C., Chen, C.-H., Hwang, F.-J., & others. (2022). Fast spatiotemporal learning framework for traffic flow forecasting. *IEEE Transactions on Intelligent Transportation Systems*.
15. Fang, Y., Zhao, F., Qin, Y., Luo, H., & Wang, C. (2022). Learning all dynamics: Traffic forecasting via locality-aware spatio-temporal joint Transformer. *IEEE Transactions on Intelligent Transportation Systems*, 23(12), 23433–23446.
16. Li, M., Tang, Y., & Ma, W. (2023). Few-sample traffic prediction with graph networks using locale as relational inductive biases. *IEEE Transactions on Intelligent Transportation Systems*, 24(2), 1894–1908.
17. Chen, Y., Li, K., Yeo, C. K., & Li, K. (2023). Traffic forecasting with graph spatial-temporal position recurrent network. *Neural Networks*, 162, 340–349.
18. Huang, Y., Song, X., Zhu, Y., Zhang, S., & others. (2023). Traffic prediction with transfer learning: A mutual information-based approach. *IEEE Transactions on Intelligent Transportation Systems*, 24(8), 8236–8252.
19. Aitha, A. R. (2024). Generative AI-Powered Fraud Detection in Workers' Compensation: A DevOps-Based Multi-Cloud Architecture Leveraging, Deep Learning, and Explainable AI. *Computer Fraud and Security*.
20. Zheng, C., Fan, X., Pan, S., Jin, H., & others. (2023). Spatio-temporal joint graph convolutional networks for traffic forecasting. *IEEE Transactions on Knowledge and Data Engineering*, 36(1), 372–385.
21. Wang, C., Zuo, K., Zhang, S., Lei, H., & others. (2023). PFNet: Large-scale traffic forecasting with progressive spatio-temporal fusion. *IEEE Transactions on Intelligent Transportation Systems*, 24(12), 14580–14597.
22. Jiang, W., & Luo, J. (2023). Graph neural network for traffic forecasting: A survey. *Expert Systems with Applications*, 207, 117921.